

Modeling Website Visits

Adrien Hitz and Robin Evans

University of Oxford

We propose a multivariate model for the number of hits on a set of popular websites, and show it to accurately reflect the behavior recorded in a data set of Internet users in the United States. We assume that the random vector of visits is distributed according to a censored multivariate normal with marginals transformed to be discrete Pareto IV and, following the ideas of Gaussian graphical models, we enforce sparsity on the inverse covariance matrix to reduce dimensionality and to visualize the dependence structure as a graph. The model allows for an easy inclusion of covariates and is useful for comprehending the behavior of Internet users as a function of their age and gender.

1 Introduction

Interest in the study of online behavior has grown with the importance of the Internet itself in our society, facilitated by the advent of rich data sets on browsing history; see for instance Bucklin and Sismeiro, 2003 and Johnson et al., 2004. Several models have been suggested for the number of hits to multiple websites or for the number of page viewed, such as a time dependent stochastic model (Park and Fader, 2004), a multivariate generalization of the negative binomial distribution (Danaher, 2007) and a Gaussian copula with negative binomial marginals (Danaher and Smith, 2011).

In this article, we consider a data set from Nielsen Holdings PLC, an information and measurement company, containing the number of visits of 19 436 users to the 99 most visited websites in the United States during one month. The table below displays hits of two users to the ten most popular websites.

Website	User 1	User 2
Google	2	1
Yahoo!	155	668
Facebook	0	0
YouTube	0	2
MSN/WindowsLive/Bing	0	0
AOL Media Network	0	0
Amazon	0	0
Ask Search Network	1	0
Wikipedia	0	0
eBay	0	0

We are treating the number of visits as i.i.d. realizations $\mathbf{x} = \{\mathbf{x}^{(k)}\}_{k=1}^n$ of a random vector $\mathbf{X} = (X_1, \dots, X_d)$ for $d = 99$ and $n = 19\,436$. Modeling these data seems challenging because the marginals X_i have very heavy-tailed distributions exhibiting peaks at zero and are strongly dependent. We will present a simple multivariate distribution that fits the data relatively well. Before starting, we leave aside about 10% of the observations and refer to them as test data.

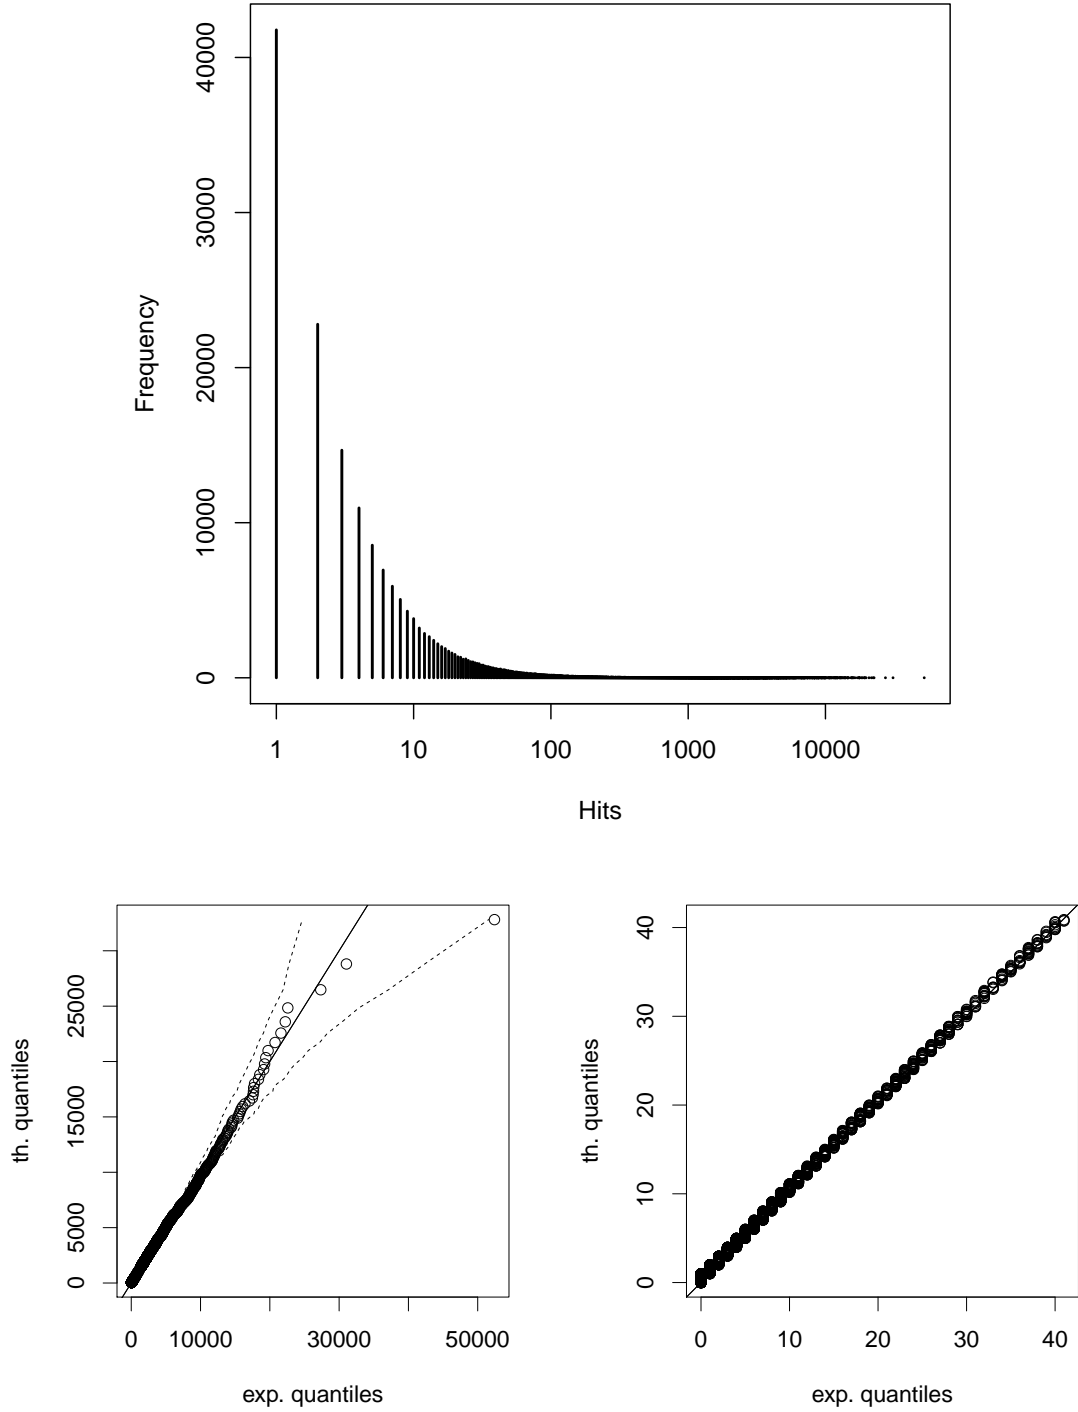


Figure 1: Frequency table of about 200 000 positive hits to any of the 99 most visited websites (top), QQ-plot for a discrete Pareto IV distribution fitted to these data (bottom left) and the same plot for quantiles smaller than 0.8 only (bottom right). The dotted lines denote 95% pointwise confidence intervals.

2 The Discrete Pareto IV Distribution

First, we seek to describe the distribution of the number of hits to any of the 99 websites, written X_{tot} . A frequency table of $X_{\text{tot}} \mid X_{\text{tot}} \geq 1$ is displayed in the top of Figure 1 on a log-scale and reveals a surprisingly regular decay. At the value zero, however, the probability mass function forms an irregular peak : 88.1% of hits are 0 while only 2.4% and 1.3% are 1 and 2 respectively. Assuming that $X_{\text{tot}} \mid X_{\text{tot}} \geq 1$ follows a Poisson or a negative binomial distribution delivers a poor fit. We thus turn our attention to a more flexible discrete family distribution called the discrete Pareto IV (D-PIV) and defined by its probability mass function $p(k) = F(k) - F(k - 1)$ on $\mathbb{N} = \{1, 2, \dots\}$, where

$$F(x) = 1 - \left\{ 1 + \xi \frac{(x + \mu)^\beta - \mu^\beta}{\sigma} \right\}^{-1/\xi} 1_{\{0 \leq x < e_1\}},$$

for $\xi \in \mathbb{R}$, $\mu \geq 0$, $\beta > 0$, $\sigma > 0$, and, if $\xi \geq 0$, $e_1 = \infty$. When $\xi > 0$, F is the cumulative distribution function (cdf) of a Pareto IV distribution (see e.g. Arnold, 2015). When $\xi < 0$, F has a finite endpoint

$$e_1 = \left(\mu^\beta + \frac{\sigma}{|\xi|} \right)^{1/\beta} - \mu,$$

and when $\xi = 0$, it is extended by continuity to $F(x) = 1 - \exp \{ (x + \mu)^\beta / \sigma - \mu^\beta / \sigma \}$. In the case $\mu = 0$, $\beta = 1$, the D-PIV coincides with a discrete generalized Pareto distribution which has been used, for instance, by Prieto et al. (2014) to model road accidents.

We fit the D-PIV to $X_{\text{tot}} \mid X_{\text{tot}} \geq 1$ by maximizing numerically the log-likelihood using the function `optim` of **R** with starting parameters $(\xi, \sigma, \beta, \mu) = (0.1, 1, 0.1, 1)$ (**R** Core Team, 2015). Even though there are more than 200 000 observations, the maximization is fast because the likelihood only needs to be evaluated once at each observed integer value. The D-PIV fits the data well as shown in the bottom of Figure 1. This is confirmed by a discrete Kolmogorov–Smirnov goodness-of-fit test in package `dgof` applied to the test data set giving a p-value of 0.14 (Arnold and Emerson, 2011). Maximum likelihood estimates and 95% confidence intervals based on asymptotic normality are displayed in the table below.

ξ	σ	β	μ
$-0.17_{\pm 0.03}$	$0.20_{\pm 0.06}$	$0.07_{\pm 0.02}$	$1.13_{\pm 0.04}$

Since $\xi < 0$, the fitted distribution has a finite endpoint e_1 , The later is roughly 1.7 times the largest observation in the data set which is 52 419 hits to Facebook.

For each website $i = 1, \dots, d$, we now fit to $X_i \mid X_i \geq 1$ a D-PIV and two embedded models corresponding to the cases $\mu = 0$ and $(\mu = 0, \beta = 1)$ and select among them according to the Bayesian information criteria (BIC). The following table shows the percentage of discrete Kolmogorov–Smirnov tests whose p-values are smaller than 0.05 and the average sample size.

	Data Set	Rejected	Av. Sample Size
D-PIV	training	5.0%	2090
D-PIV	test	6.1%	234

The rejection rate is close to 5%, indicating a good fit which supports the adequacy of the D-PIV family distribution to model $X_i \mid X_i \geq 1$. The saturated model and the embedded models in the case $\mu = 0$ and $(\mu = 0, \beta = 1)$ were selected 25, 59 and 15 times respectively.

3 The Censored Student Copula Graphical Model

We are now interested in modeling the dependence structure of $\mathbf{X} = (X_1, \dots, X_d)$. The problem cannot easily be reduced to lower dimension by assuming independence between blocks of marginals because the dependence between them is rather strong. To illustrate this, we define a binary vector \mathbf{B} by $B_i = 1_{\{X_i \geq 1\}}$, where 1_A is the indicator function, and perform χ^2 tests of independence with level 0.05 for all pairs; the test rejects independence 98% of the time. Graphical models provide a way to simplify a joint density by making assumptions of conditional independence between marginals as we will discuss later. We use `ci.test` of package `bnlearn` to test conditional independence $B_i \perp\!\!\!\perp B_j \mid \mathbf{B}_{-i,-j}$ (Scutari, 2010). Pairwise conditional independence is rejected only 33% of the time, suggesting that there might be many conditional independence relations to exploit to reduce dimensionality.

Is there a simple multivariate distribution that can be chosen for \mathbf{X} ? There are extensions of the Pareto IV in the multivariate case but they typically have limited dependence structure or are intractable (Arnold, 2015). As explained by copula theory, the marginals

of any multivariate continuous distribution can be transformed to obtain another distribution with marginals of arbitrary continuous distribution while preserving the dependence structure (see e.g. Sklar, 2010 or Nelsen, 2006). In the discrete case, copulas suffer from limitations such as non-identifiability (Genest and Nešlehová, 2007).

An additional difficulty for modeling \mathbf{X} , as we have seen, is that its distribution is fundamentally different on zero than on positive integers, requiring a discrete multivariate distribution that accounts for this mixture. The multivariate zero-inflated Poisson distribution is such an instance (Li et al., 1999; Liu and Tian, 2015).

The multivariate Gaussian and Student distributions are some of the few known multivariate distributions with explicit multivariate marginal and conditional densities. Their dependence structure is well understood and has some flexibility without being overdetermined as it only involves pairwise interactions. We now explain how they can be used for modeling \mathbf{X} . Let $\mathbf{t} \in \mathbb{R}^d$ be a vector of thresholds and suppose that \mathbf{Z} follows a centered Student distribution with degree of freedom $\nu \in (2, \infty]$ and covariance matrix $\frac{\nu}{\nu-2}\Sigma$ such that Σ has diagonal $\mathbf{1} \in \mathbb{R}^d$. By convention, $\mathbf{Z} \sim \mathcal{N}(0, \Sigma)$ in the case $\nu = \infty$. We obtain a vector \mathbf{X} with values in $(\{0\} \times \mathbb{N}_+)^d$ satisfying $X_i \mid X_i \geq 1 \sim F_{\text{D-PIV}}$ for all i by applying to \mathbf{Z} the following procedure: censor Z_i when it falls below t_i , set censored values to 0, transform non-censored Z_i appropriately and round them. More precisely, define \mathbf{X} by

$$X_i = \begin{cases} \lfloor Y_i \rfloor + 1 & \text{if } Z_i \geq t_i, \\ 0 & \text{else,} \end{cases} \quad Y_i := F_{\text{PIV}_i}^{-1}\{\Phi_{Z|Z \geq t_i}(Z_i \mid Z_i \geq t_i)\},$$

for all $i = 1, \dots, d$, where $\Phi_{Z|Z \geq t_i}$ is the cdf of a truncated Student and F_{PIV_i} is such that $F_{\text{PIV}_i}(k) - F_{\text{PIV}_i}(k-1)$ is the probability mass function of the D-PIV fitted to $X_i \mid X_i \geq 1$. Here, $\lfloor x \rfloor$ denotes the largest integer smaller or equal to x .

This formulation provides a simple yet non trivial probability distribution with values in \mathbb{N}^d accounting for the irregularity of the distribution between null and positive values. It has been frequently used to model multivariate rainfall data where zeros occur when no rain is measured, see e.g. Bell (1987) or Allcroft and Glasbey (2003). It also appears in multivariate extreme value analysis where zeros correspond to non-extremal events, such as a modeling of extremal oceanographic data in Bortot et al. (2000).

In theory, we could express the probability mass function of \mathbf{X} as a sum of terms involving the joint cdf of \mathbf{Z} , but the expression is intractable when d is large. We are thus treating positive values of \mathbf{X} as continuous, making the working assumption that $X_i \approx Y_i + \frac{1}{2}$ when $Z_i \geq t_i$. A similar approximation for the multivariate normal copula in the case of discrete data is studied in Nikoloulopoulos (2016). This is a reasonable assumption here because the range of positive values is relatively large. Another benefit is that it yields a tractable expression for conditional distributions: for any disjoint sets $A, B, C, D \subseteq \{1, \dots, d\}$, $\mathbf{k} \in \mathbb{N}_+^d$,

$$\Pr\{\mathbf{X}_{AB} = (\mathbf{k}_A, \mathbf{0}) \mid \mathbf{X}_{CD} = (\mathbf{k}_C, \mathbf{0})\} = \frac{\Phi(\mathbf{t}_{BD} \mid \mathbf{Z}_{AC} = \mathbf{z}_{AC})}{\Phi(\mathbf{t}_B \mid \mathbf{Z}_{AC} = \mathbf{z}_{AC})} \phi(\mathbf{z}_A \mid \mathbf{z}_C), \quad (1)$$

where $z_i = \Phi_{Z_i|Z_i \geq t_i}^{-1}\{F_{\text{PIV}}(\mathbf{k}_i - \frac{1}{2})\}$. The quantities $\Phi(\cdot \mid \mathbf{z}_F)$ and $\phi(\cdot \mid \mathbf{z}_F)$ denotes respectively the joint cdf and density function of $\mathbf{Z}_E \mid \mathbf{Z}_F = \mathbf{z}_F$ for some disjoint sets $F, E \subset \{1, \dots, d\}$ and coincides with the joint cdf $\Phi(\cdot)$ and density $\phi(\cdot)$ of a non-centered Student distribution on $\mathbb{R}^{|E|}$ (Ding, 2016). The joint cdf $\Phi(\cdot)$ can be evaluated using the package `mvtnorm` (Genz et al., 2008). Notice that one can sample from $\mathbf{Z}_E \mid \mathbf{Z}_F$ using (1), for instance, by applying inverse transform sampling recursively.

We refer to the model above as the censored Student copula model and we now present a possible way of estimating it. Recall that the website visits $\mathbf{x} = \{\mathbf{x}^{(k)}\}_{k=1}^n$ are assumed to be sampled from \mathbf{X} . The parameters t_i can directly be estimated as $\hat{t}_i = \Phi^{-1}(n_i/n)$ where n_i is the number of observations such that $x_i^{(k)} = 0$. It remains to estimate the matrix Σ to fully determine the distribution. Its maximum likelihood estimate in the Gaussian case would be

$$\hat{S} = \frac{1}{n} \sum_{k=1}^n \mathbf{z}^{(k)} \mathbf{z}^{(k)T},$$

if a sample \mathbf{z} from \mathbf{Z} was observed. However, only a censored transformation of \mathbf{z} is observed and different methods are possible in this case (see e.g. Lee and Scott, 2012 or Schemper et al., 2013). We found satisfactory performance and efficiency of estimating $\rho = \Sigma_{ij}$ by maximizing the pairwise likelihood $\ell_{ij}(\rho) = \sum_{k=1}^n \log \Pr(X_i = x_i^{(k)}, X_j = x_j^{(k)})$

for all pairs (i, j) separately, where

$$\begin{aligned}\Pr(X_i = 0, X_j = 0) &= \Pr(Z_i \leq t_i, Z_j \leq t_j), \\ \Pr(X_i = x_i, X_j = x_j) &= c_i(x_i)c_j(x_j)\Pr(Z_i = z_i, Z_j = z_j), \\ \Pr(X_i = x_i, X_j = 0) &= c_i(x_i)\Pr(Z_i = z_i, Z_j \leq t_j), \\ z_i &= M_i^{-1}\{F_{X_i|X_i \geq 1}(x_i)\}, \quad M_i(z) = \Pr(Z_i \leq z \mid Z_i \geq t_i), \quad c_i(x_i) = \frac{f_{X_i|X_i \geq 1}(x_i)}{M_i'(z_i)},\end{aligned}$$

where we abuse notation and write $\Pr(Z_i = z) := \frac{\partial}{\partial z}\Pr(Z_i \leq z)$. All together, this gives a consistent estimator $\hat{\Sigma}^C$ of Σ defined by $\hat{\Sigma}_{ij}^C = \arg \max_{\rho \in [-1, 1]} \ell_{ij}(\rho)$.

We compute Σ^C in the case $\nu = \infty$ for $n = 19436$ users and display the largest positive and negative correlations Σ_{ij}^C in the table below with a 95% confidence interval based on asymptotic normality of the maximum likelihood estimator.

i	j	ρ_{ij}
TurboTax	Intuit	$0.91_{\pm 0.01}$
YouTube	YouTube-NoCookie	$0.86_{\pm 0.00}$
YouTube	Vevo	$0.78_{\pm 0.01}$
Expedia	TripAdvisor	$0.75_{\pm 0.02}$
YouTube-NoCookie	Vevo	$0.74_{\pm 0.01}$
Lowe's	The Home Depot	$0.71_{\pm 0.02}$
eBay	PayPal	$0.68_{\pm 0.02}$
Facebook	Zynga	$0.68_{\pm 0.02}$
Blogger	WordPress	$0.67_{\pm 0.02}$
U.S. Internal Revenue Service	TurboTax	$0.65_{\pm 0.03}$
American Express	Vevo	$-0.11_{\pm 0.06}$
Comcast Digital Entertainment	Road Runner	$-0.10_{\pm 0.05}$
Vevo	Discover	$-0.08_{\pm 0.06}$

The marked dependence between the websites above is sensible: Turbo Tax and Intuit are tax softwares, YouTube and Vevo video websites, Expedia and TripAdvisor travel websites, Lowe's and The Home Depot home improvement stores, eBay is an e-commerce company who owned PayPal, an online payment operator. Facebook, a social networking service, had a partnership with Zynga, an online video game provider; Blogger and WordPress are blog publishing services.

Although there are many other positively correlated pairs of websites, only 3 pairs exhibit a significant negative correlation according to our model. Visitors of Vevo are less interested in American Express and Discover, two financial service companies, as are visitors of Road Runner, a running shoe store, in Comcast, a mass media company.

Two characteristics are available on users: their age, which takes tabulated values and is censored above 65, and their gender, which is binary. These covariates could be included in the analysis by regressing the parameters of the model with respect to them. We prefer here to treat them on the same terms as the other variables in a fully multivariate analysis. Let Z_a and Z_b be Student distributed. For age, we assume that $X_a = 1_{\{Z_a \geq t_a\}}$ for some $t_a \in \mathbb{R}$. For gender, we suppose that X_b is censored from above when $Z_b \geq t_b$ for some t_b and $X_b = F_b^{-1}\{\Phi_{Z_b|Z_b \leq t_b}(Z_b | Z_b \leq t_b)\}$, where F_b^{-1} is the empirical quantile function of \mathbf{x}_b . Note that since \mathbf{x}_b is discrete but treated as continuous, it is a better approximation to transform it to uniform using $\{F_b(x) + F_b(x - 1)\}/2$ instead of its empirical cdf F_b . Following the procedure explained previously, we obtain an estimate for the threshold vector $\mathbf{t} \in \mathbb{R}^{d+2}$ and the matrix $\Sigma \in \mathbb{R}^{d+2} \times \mathbb{R}^{d+2}$. The underlying random vector is now $\mathbf{Z} = (Z_1, \dots, Z_{99}, Z_a, Z_b)$ which is Student distributed with degree of freedom $\nu > 2$, mean $\mathbf{0}$ and covariance matrix $\frac{\nu}{\nu-2}\Sigma$.

The next table shows with which websites age and gender share the strongest correlation in the case $\nu = \infty$, with the convention $X_a = 1$ for female and $X_a = 0$ for male.

i	j	Σ_{ij}^C	i	j	Σ_{ij}^C
age	Legacy	$0.34_{\pm 0.03}$	gender	Pinterest	$0.31_{\pm 0.04}$
age	American Express	$0.29_{\pm 0.04}$	gender	JCPenney	$0.30_{\pm 0.04}$
age	WhitePages	$0.28_{\pm 0.03}$	gender	Macy's	$0.29_{\pm 0.04}$
age	Shopzilla	$0.27_{\pm 0.03}$	gender	Allrecipes	$0.28_{\pm 0.03}$
age	SuperPages	$0.27_{\pm 0.03}$	gender	Everyday Health	$0.27_{\pm 0.03}$
age	Vevo	$-0.30_{\pm 0.02}$	gender	ESPN	$-0.25_{\pm 0.03}$
age	YouTube-NoCookie	$-0.23_{\pm 0.02}$	gender	Turner-SI	$-0.22_{\pm 0.03}$
age	Disney DOL	$-0.22_{\pm 0.03}$	gender	Big Lead Sports	$-0.21_{\pm 0.04}$
age	YouTube	$-0.21_{\pm 0.02}$	gender	NFL	$-0.21_{\pm 0.03}$
age	Tumblr	$-0.20_{\pm 0.03}$	gender	CNET	$-0.17_{\pm 0.04}$

The most visited websites by older persons were Legacy, an online memorial provider, American Express, a financial service company, White Pages, a contact information provider, Shopzilla, a shopping website and Experian, a global information services group. Younger persons mostly hit to Vevo, YouTube, Disney, a mass media and entertainment company, and Tumblr, a social networking website.

Women were more frequently on Pinterest, a photo sharing website, JCPenney and Macy's, two department store chains, Allrecipes, a social networking service focused on food, and Everyday Health, a company producing content on wellness. Men, on the other hand, preferred sport related networks and CNET, a technology news website.

From a model selection point of view, it is natural to try to represent the matrix Σ with fewer parameters. As we have seen, the marginals are strongly dependent so Σ is unlikely to be sparse. Following the ideas of graphical models, one can try to exploit sparsity of the inverse matrix Σ^{-1} which translates, in the Gaussian case, into conditional independence relations between some of the marginals: $\Sigma_{ij}^{-1} = 0$ if and only if $Z_i \perp\!\!\!\perp Z_j \mid \mathbf{Z}_{\{1,\dots,d\}\setminus\{i,j\}}$ (see e.g. Lauritzen, 1996). The constraint is a little different for the Student distribution: $\Sigma_{ij}^{-1} = 0$ implies that Z_i and Z_j are conditionally uncorrelated given the rest of the vector (Finegold and Drton, 2011). Besides reducing the number of parameters, these assumptions allow us to visualize the dependence structure as a graph $\mathcal{G} = (V, E)$ defined as follows: its set of nodes V is $\{1, \dots, d\}$, i.e., each node corresponds to a marginal of \mathbf{Z} , and its set of edges E satisfies

$$(i, j) \notin E \implies \Sigma_{ij}^{-1} = 0. \quad (2)$$

An efficient procedure to estimate a sparse matrix Σ^{-1} given a sample \mathbf{z} i.i.d. drawn from $\mathcal{N}(0, \Sigma)$ is proposed in Friedman et al. (2008). It solves the convex optimization problem called Gaussian graphical lasso which consists in maximizing the log-likelihood of a multivariate Gaussian with an additional penalization to enforce sparsity. More precisely,

$$\hat{\Theta} = \arg \max_{\Theta \succeq 0} \log \det \Theta - \text{tr}(\hat{S}_n \Theta) - \lambda_n \|\Theta\|_1, \quad (3)$$

where $\hat{S}_n = n^{-1} \mathbf{z}^T \mathbf{z}$ is the empirical covariance matrix, $\lambda_n > 0$ is a regularization parameter, “tr” denotes the trace of a matrix, “ $\succeq 0$ ” means non-negative definite and

$\|\Theta\|_1 := \sum_{ij} |\Theta_{ij}|$. After slightly reformulating (3), Ravikumar et al. (2011) show that for any consistent estimator \hat{S}_n of Σ (such as $\hat{\Sigma}^C$) and under an additional assumption (called incoherence or irrepresentability condition), the solution $\hat{\Theta}$ is a consistent estimator of Σ^{-1} and it correctly detects null entries as $n \rightarrow \infty$ for fixed d . When $\nu < \infty$, solving (3) still makes sense as $\hat{\Theta}$ is the closest matrix to Σ^{-1} in terms of a Bregman divergence. Note that if one is interested in estimating the graph \mathcal{G} only, Meinshausen and Bühlmann (2006) present a procedure to estimate it consistently in the Gaussian case. There is a rich literature on Gaussian graphical lasso (see e.g. Banerjee et al., 2008; Yuan, 2010; Cai et al., 2016), copula Gaussian graphical models (Dobra and Lenkoski, 2011; Xue and Zou, 2012; Liu et al., 2012), Student graphical lasso (Finegold and Drton, 2011) and Gaussian lasso in the censored case (Johnson, 2009), including several variants such as the adaptive lasso which penalizes coefficients differently (Zou, 2006) and a decomposition of Σ^{-1} into the sum of a sparse and a low-rank matrix (Chandrasekaran et al., 2012).

Coming back to our data analysis, we compute $\hat{S}_n = \hat{\Sigma}^C$ for various degrees of freedom ν and solve (3) for several regularization parameters λ using the package **huge** (Zhao et al., 2012) to obtain matrices $\hat{\Sigma}$ whose inverse are sparse. We then choose ν and λ by minimizing two scores: the average negative log-likelihood

$$\ell(\hat{\Sigma}) = -\frac{1}{n} \sum_{k=1}^n \left\{ \log \Phi \left(\mathbf{t}_{A_k^c} \mid \mathbf{z}_{A_k}^{(k)} \right) + \log \phi \left(\mathbf{z}_{A_k}^{(k)} \right) + \sum_{i \in A_k} \log c_i \left(x_i^{(k)} \right) \right\}, \quad (4)$$

where $A_k = \{i : x_i^{(k)} \neq 0\}$, and the sum of average pairwise negative log-likelihood

$$\ell_{\text{pairwise}}(\hat{\Sigma}) = \sum_{i < j} \ell_{ij}(\hat{\Sigma}_{ij}).$$

We compute ℓ and ℓ_{pairwise} from the test data set using respectively 400 and 19 436 observations and display them in the table below. We also report the percentage of null entries in the upper triangular part of $\hat{\Sigma}^{-1}$ referred to as sparsity.

		$\hat{\Sigma}^C$	$\hat{\Sigma}_{\lambda=0.0075}$	$\hat{\Sigma}_{\lambda=0.02}$	$\hat{\Sigma}_{\lambda=0.1}$
$\nu = \infty$	sparsity	0%	25%	50%	73%
	ℓ_{pairwise}	3079.33	3079.31	3079.36	3081.78
	ℓ	23.14	23.07	23.09	23.43
$\nu = 60$	sparsity	0%	24%	50%	75%
	ℓ_{pairwise}	3079.79	3079.78	3079.83	3082.15
	ℓ	23.20	23.12	23.11	23.43
$\nu = 30$	sparsity	0%			
	ℓ_{pairwise}	3080.19			
	ℓ	23.24			

Both scores agree on selecting the censored Gaussian copula graphical model with $\lambda = 0.0075$, corresponding to a sparsity of 25%. Figure 2 shows the corresponding graph, plotting only edges (i, j) such that $\hat{\Sigma}_{ij}^{-1}$ are among the 5% largest non-null entries in absolute values. An edge between i and j should be interpreted as X_i carrying relevant information to predict X_j . We identify several meaningful clusters such as financial companies (Chase, Discover, Citibank, Capital One), news (Fox News, CNN, ABCNEWS, USATODAY, NYTimes, Tribune Newspaper) and sports (Turner-SI, NFL Internet Network, Big Lead Sports, ESPN). The graph also illustrates which websites a user tends to visit depending on its age and gender and we recognize some of the relations discussed earlier.

We now present several diagnostics of the selected model. We start by testing if pairs from the real data set have the probability expected from the model of occurring in certain regions. Binomial tests with level 0.05 are performed for every pairs (i, j) and the result is reported in the next table.

Region	Rejected
$\Pr(X_i > 0, X_j > 0)$	10.4%
$\Pr(X_i = 0, X_j > 0)$	7.7%

¹The following websites do not have any neighbor in the graph in Figure 2: Apple, Glam Media, Weather Channel Network, Bank of America, Adobe, Wells Fargo, NBC Universal Sites, Napster, ShopAtHome, Hewlett Packard, Experian.

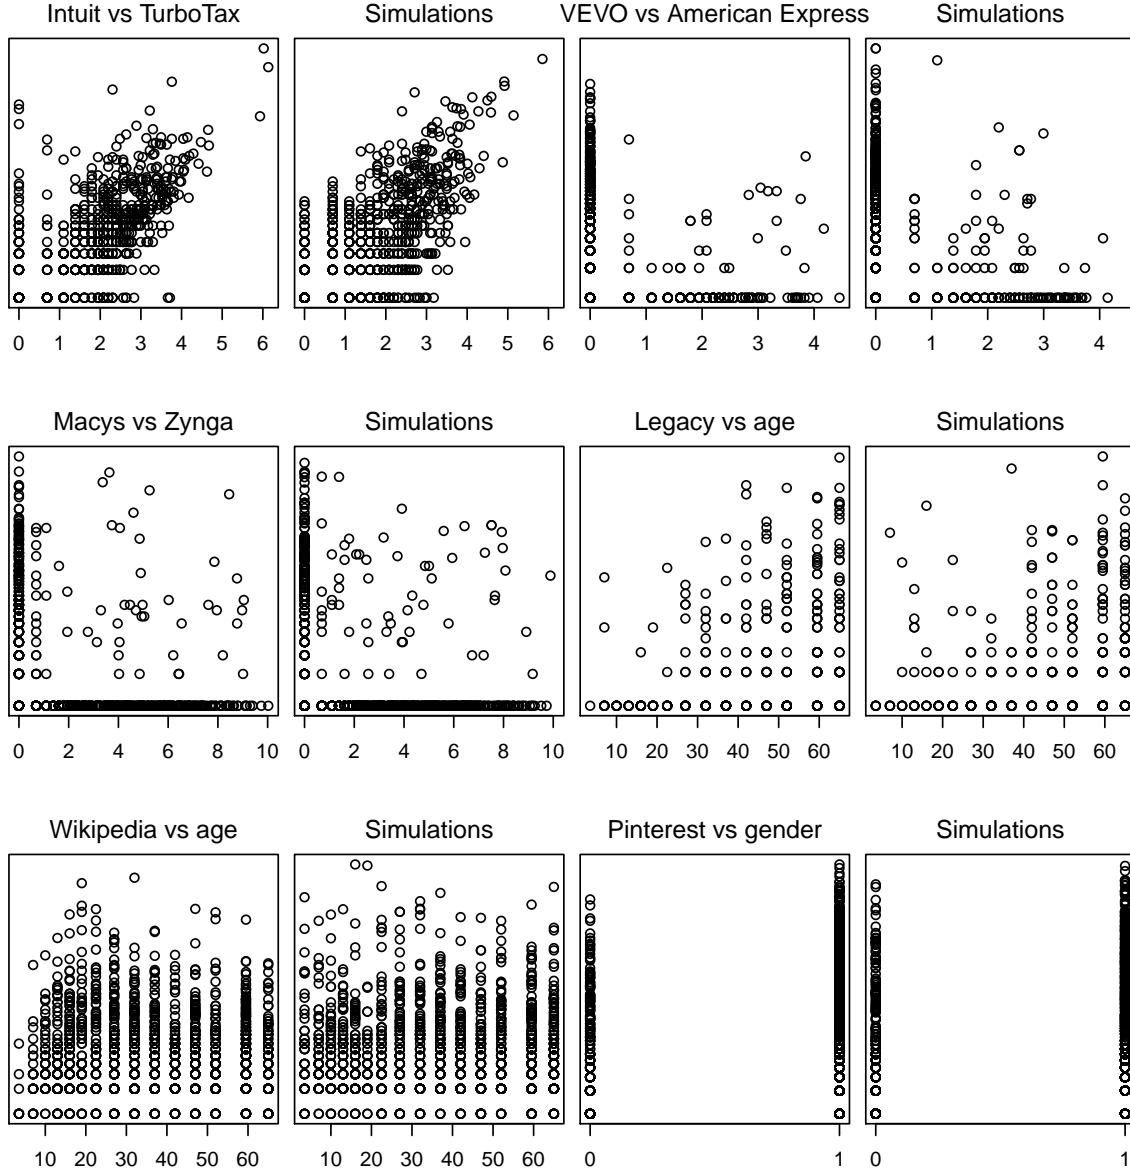


Figure 3: Several scatterplots of the data (first and third columns) and simulations from a censored Gaussian copula graphical model (second and fourth columns). The quantities plotted are the age X_a , the gender X_b and a transformation $\log(X_i + 1)$ of the number of hits X_i to website i .




Figure 4: Diagnostic QQ-plots for the sum of hits to the 5 (left) and 99 (right) most visited websites. Empirical quantiles of the data are plotted against simulated data from the censored Gaussian copula graphical model on a log scale.

The percentage of rejection is above 5%, the rate of a correct model.

As a further diagnostic, we draw 19 436 realizations from the censored Gaussian copula graphical model and compare them to the data. Figure 3 displays a few scatterplots chosen as follows: most positively, most negatively and least correlated pairs of website hits; most positively and less correlated pairs involving age; most correlated pair involving gender. Replicates from the model appear to be relatively similar to the data. As another comparison, Figure 4 shows QQ-plots for the sum of hits to the 5 and 99 most visited websites. The quantiles form a slightly concave line instead of a straight line, revealing some inaccuracies.

We end by assessing the performance of the model throughout the following experiment. We try to predict if $X_i > 0$ given \mathbf{X}_{-i} . For $i = 1, \dots, 99$, this corresponds to predicting if a user visits website i given its age, gender and visits on the other websites during the month. For gender, it corresponds to predicting if the user is a female. For age, we try to predict if $X_a > 35$, i.e., if the user is older than 35. To gain efficiency, we make the working assumption that $X_i \perp\!\!\!\perp X_j \mid \mathbf{X}_{\{1, \dots, d\} \setminus \{i, j\}}$ if $\Sigma_{ij}^{-1} = 0$, a relation which holds for \mathbf{Z} but only

approximately for \mathbf{X} . We choose the score function

$$1_{\{X_i > 0\}} \log \hat{p} + (1 - 1_{\{X_i > 0\}}) \log(1 - \hat{p}),$$

where \hat{p} is the predicted probability that $X_i > 0$, and compare the predictions of the model to various methods. First, we naively approximate $\Pr(X_i > 0 \mid \mathbf{X}_{-i})$ by $\Pr(X_i > 0)$. Second, we fit an Ising model — a probabilistic graphical model for binary data — to $\mathbf{B} = (1_{X_1 > 0}, \dots, 1_{X_{d+2} > 0})$ using the R package `IsingFit` and compute $\Pr(X_i > 0 \mid \mathbf{B}_{-i})$ (Van Borkulo and Epskamp, 2014). Third, we train decision tree algorithms found in packages `rpart`, `tree`, `ctree` and `randomForest` (using 50 trees for the latter) to predict $X_i > 0 \mid \mathbf{X}_{-i}$, (Van Borkulo and Epskamp, 2014; Therneau et al., 2015; Hothorn et al., 2006; Liaw and Wiener, 2002). The next table shows the average score and the percentage of correct guessing computed for about 400 predictions from the test data set.

Naive	$\hat{\Sigma}^C$	$\hat{\Sigma}_{\lambda=0.0075}$	$\hat{\Sigma}_{\lambda=0.02}$	$\hat{\Sigma}_{\lambda=0.1}$	Ising	<code>rpart</code>	<code>tree</code>	<code>ctree</code>	RF
0.334	0.254	0.255	0.254	0.257	0.267	0.324	0.293	0.287	∞^2
86.8%	89.6%	89.6%	89.3%	89.9%	89.3%	88.5%	88.5%	88.1%	89.8%

The censored Gaussian copula is a simple probabilistic model that provides an approximation of the full joint distribution of the data and is thus valuable to comprehend its dependence structure as a whole. The experiment above suggests that decision trees that were specifically trained for certain predictions do not outperform the model, at least when they are used without further tuning. The Ising model performs relatively well and seems a suitable alternative if only the fact that a website is visited matters for the analysis, but not the number of hits. On average, the censored Gaussian gives the best performance, showing its ability to capture relevant information in this data set.

Acknowledgements We would like to thank Nielsen Holdings for providing the data, and the first author is grateful to the Berrow Foundation for financial support.

²`RandomForest` returns 7 infinite values; its average score computed from finite values only is 0.221.

References

- D. J. Allcroft and C. A. Glasbey. A latent Gaussian Markov random-field model for spatiotemporal rainfall disaggregation. *Journal of the Royal Statistical Society. Series C. Applied Statistics*, 52(4):487–498, 2003.
- B. C. Arnold. *Pareto distribution*. Wiley Online Library, 2015.
- T. B. Arnold and J. W. Emerson. Nonparametric goodness-of-fit tests for discrete null distributions. *The R Journal*, 3(2):34–39, 2011.
- O. Banerjee, L. E. Ghaoui, and A. d’Aspremont. Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *Journal of Machine Learning Research*, 9:485–516, 2008.
- T. L. Bell. A space-time stochastic model of rainfall for satellite remote-sensing studies. *Journal of Geophysical Research: Atmospheres*, 92(D8):9631–9643, 1987.
- P. Bortot, S. G. Coles, and J. Tawn. The multivariate Gaussian tail model: an application to oceanographic data. *Journal of the Royal Statistical Society. Series C. Applied Statistics*, 49(1):31–49, 2000.
- R. E. Bucklin and C. Sismeiro. A model of web site browsing behavior estimated on clickstream data. *Journal of marketing research*, 40(3):249–267, 2003.
- T. T. Cai, W. Liu, and H. H. Zhou. Estimating sparse precision matrix: Optimal rates of convergence and adaptive estimation. *The Annals of Statistics*, 44(2):455–488, 2016.
- V. Chandrasekaran, P. A. Parrilo, and A. S. Willsky. Latent variable graphical model selection via convex optimization. *The Annals of Statistics*, 40(4):2005–2013, 2012.
- P. J. Danaher. Modeling page views across multiple websites with an application to internet reach and frequency prediction. *Marketing Science*, 26(3):422–437, 2007.
- P. J. Danaher and M. S. Smith. Modeling multivariate distributions using copulas: applications in marketing. *Marketing Science*, 30(1):4–21, 2011.

- P. Ding. On the conditional distribution of the multivariate t distribution. *The American Statistician*, 70(3):293–295, 2016.
- A. Dobra and A. Lenkoski. Copula Gaussian graphical models and their application to modeling functional disability data. *The Annals of Applied Statistics*, 5(2A):969–993, 2011.
- M. Finegold and M. Drton. Robust graphical modeling of gene networks using classical and alternative t -distributions. *The Annals of Applied Statistics*, 5(2A):1057–1080, 2011.
- J. Friedman, T. Hastie, and R. Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.
- C. Genest and J. Nešlehová. A primer on copulas for count data. *Astin Bulletin*, 37(02):475–515, 2007.
- A. Genz, F. Bretz, T. Miwa, X. Mi, F. Leisch, F. Scheipl, and T. Hothorn. R package mvtnorm: multivariate normal and t distributions, 2008.
- T. Hothorn, K. Hornik, and A. Zeileis. Unbiased recursive partitioning: A conditional inference framework. *Journal of Computational and Graphical statistics*, 15(3):651–674, 2006.
- B. A. Johnson. On lasso for censored data. *Electronic Journal of Statistics*, 3:485–506, 2009.
- E. J. Johnson, W. W. Moe, P. S. Fader, S. Bellman, and G. L. Lohse. On the depth and dynamics of online search behavior. *Management Science*, 50(3):299–308, 2004.
- S. L. Lauritzen. *Graphical Models*. Oxford University Press, 1996.
- G. Lee and C. Scott. EM algorithms for multivariate Gaussian mixture models with truncated and censored data. *Computational Statistics & Data Analysis*, 56(9):2816–2829, 2012.
- C.-S. Li, J.-C. Lu, J. Park, K. Kim, P. A. Brinkley, and J. P. Peterson. Multivariate zero-inflated Poisson models and their applications. *Technometrics*, 41(1):29–38, 1999.

- A. Liaw and M. Wiener. Classification and regression by randomforest. *R news*, 2(3): 18–22, 2002.
- H. Liu, F. Han, M. Yuan, J. Lafferty, and L. Wasserman. High-dimensional semiparametric Gaussian copula graphical models. *The Annals of Statistics*, 40(4):2293–2326, 2012.
- Y. Liu and G.-L. Tian. Type I multivariate zero-inflated Poisson distribution with applications. *Computational Statistics & Data Analysis*, 83:200–222, 2015.
- N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*, 34(3):1436–1462, 2006.
- R. B. Nelsen. *An introduction to copulas*. Springer Series in Statistics. Springer, second edition, 2006.
- A. K. Nikoloulopoulos. Efficient estimation of high-dimensional multivariate normal copula models with discrete spatial responses. *Stochastic Environmental Research and Risk Assessment*, 30(2):493–505, 2016.
- Y.-H. Park and P. S. Fader. Modeling browsing behavior at multiple websites. *Marketing Science*, 23(3):280–303, 2004.
- F. Prieto, E. Gómez-Déniz, and J. M. Sarabia. Modelling road accident blackspots data with the discrete generalized Pareto distribution. *Accident Analysis & Prevention*, 71: 38–49, 2014.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, 2015.
- P. Ravikumar, M. J. Wainwright, G. Raskutti, and B. Yu. High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935–980, 2011.
- M. Schemper, A. Kaider, S. Wakounig, and G. Heinze. Estimating the correlation of bivariate failure times under censoring. *Statistics in Medicine*, 32(27):4781–4790, 2013.

- M. Scutari. Learning Bayesian networks with the bnlearn R package. *Journal of Statistical Software*, 35(3):1–22, 2010.
- A. Sklar. Fonctions de répartition a n dimensions et leurs marges. *Annales de l'I.S.U.P.*, 54(1-2):3–6, 2010.
- T. Therneau, B. Atkinson, and B. Ripley. R package rpart: Recursive partitioning for classification, regression and survival trees, 2015.
- C. Van Borkulo and S. Epskamp. R package IsingFit: Fitting Ising models using the eLasso method, 2014.
- L. Xue and H. Zou. Regularized rank-based estimation of high-dimensional nonparanormal graphical models. *The Annals of Statistics*, 40(5):2541–2571, 2012.
- M. Yuan. High dimensional inverse covariance matrix estimation via linear programming. *Journal of Machine Learning Research*, 11:2261–2286, 2010.
- T. Zhao, H. Liu, K. Roeder, J. Lafferty, and L. Wasserman. The **huge** package for high-dimensional undirected graph estimation in R. *Journal of Machine Learning Research*, 13:1059–1062, 2012.
- H. Zou. The adaptive lasso and its oracle properties. *Journal of the American Statistical Association*, 101(476):1418–1429, 2006.